Appendix C. Upgrade to 4 Megawatts (Phase II)

Chuck Ankenbrandt

A note from the author: This appendix describes an upgrade path to a facility having a beam power of 4 MW to support a muon collider. This upgrade is called Phase II of the Proton Driver and is based upon a 1 GeV linac and a 3 GeV Pre-Booster. As I began to think about how to write this appendix, I realized that I had written a similar document, "Design Concepts for Fermilab Proton Source Rings", for the 1997 Fermilab Proton Source Summer Study. (That document appeared in Fermilab Technical Memo TM-2021, "A Development Plan for the Fermilab Proton Source", edited by Steve Holmes.) It would have been expedient to simply edit my original document to reflect the design changes that have occurred in the intervening three years. However, after further thought I decided instead that it would be more informative, more interesting, perhaps even more educational for student readers (albeit perhaps also more self-indulgent) to produce this appendix by annotating that document to indicate how our thinking has evolved and where design changes have been made since it was written. Accordingly, the text of that document is reproduced here in normal typeface, and my notes about subsequent design developments appear in italics.

Design Concepts for Fermilab Proton Source Rings

Chuck Ankenbrandt
August 4, 1997

C.1. Introduction

The Fermilab physics programs of the future need a reliable, high-performance proton source. A straightforward approach to meet the diverse and demanding needs of those programs is described. In particular, the considerations that led to the choice of first-iteration values for major parameters for the synchrotron rings are presented.

The perfect proton source for Fermilab would be able to deliver beams having the ideal beam parameters for all possible future physics programs. Among the possibilities presently envisioned are beams for the Tevatron collider and for the VLHC, for antiproton production, for fixed-target physics based on the Main Injector, for experiments such as miniBOONE that use beam from the source directly, and for muon production for a muon collider.

The muon collider makes such severe demands on the proton source that it tends to dominate design considerations. Over the last few years, various approaches to meet those needs have been explored. Within the last year a simple approach has been developed that not only satisfies the requirements of the muon collider but also can be adjusted, by appropriate parameter choices, to match the needs of the rest of the future program.

The main body of this design study report describes a facility intended to support the Phase I applications. (Note that a neutrino factory based on a muon storage ring was apparently not yet on my radar screen in 1997.) The muon collider originally drove the development of the concepts described in this appendix and would presumably be the main customer of the Phase II upgrades.

C.2. The Muon Collider Requirements

Two short (σ = 2 nsec) bunches each containing 5×10^{13} protons at an energy around 16 GeV and a repetition rate of 5 Hz would meet the needs of a high-performance muon collider. (The proton source requirements are about the same over the whole range of final muon collider energies that have been considered, from 50×50 GeV up to 2×2 TeV.) These specifications deserve some elaboration.

Subsequently the muon collider designers requested four proton bunches in order to reduce beam loading effects during muon acceleration and beam-beam effects in the collider. Accordingly, the present scenario envisages 4 bunches each containing 2.5×10^{13} protons.

A kinetic energy of 8 GeV rather than a higher energy would seem a natural choice for a proton source at Fermilab; 8 GeV is the energy of the present Booster and antiproton source as well as the design injection energy of the Main Injector. However, the performance of the Main Injector is likely to benefit from raising its injection energy. Not only would space-charge effects be alleviated, but the available normalized aperture would roughly double, scaling with momentum, and the injection field quality would improve. (The ultimate benefit would derive from raising the injection energy above the Main Injector transition energy of about 20 GeV, or alternatively lowering the transition energy.) Furthermore, the dependence of pion production on proton energy is thought to be such that almost twice as many protons would be needed at 8 GeV as at 16 GeV. (There are significant uncertainties in the pion production cross sections in normalization, distributions, and energy dependence. There are disagreements amongst the hadroproduction model programs, and the data are insufficient to resolve the discrepancies. There is an experiment in progress at the Brookhaven AGS intended to shed light on the situation.)

The present design has an output kinetic energy of 16 GeV; that value results from constraining the circumference of the second ring to match that of the existing Booster. There are two main advantages of this choice. First, the second ring could occupy the same tunnel as the relocated Booster; secondly, the beam batch length from the second ring would match that of the present Booster, which would be ideal for antiproton production. Note, however, that the existing Booster magnets can not go beyond about 10 GeV.

The design choices described in these two paragraphs have evolved considerably, but the issues discussed remain controversial. First of all, more work on measuring and

modeling pion production and muon capture has clarified the situation considerably, confirming the rough scaling with beam power implied above. (Deviations from the scaling rule of thumb favor lower beam energies, particularly for carbon targets.) Regarding the choice of accelerator parameters, the lattice designers were unable to achieve all the desired features of the lattice, including separate function and transition avoidance, in a 16-GeV machine having the same circumference as the Booster. The decision was made to increase the circumference to 1.5 times that of the Booster rather than lowering the maximum energy from 16 GeV. However, the purported benefits of increasing the energy of injection into the Main Injector seem marginal at best unless it is possible to get above transition in the Main Injector, which would be very expensive at 15 Hz. While it is true that the Main Injector would presumably accept a beam of larger normalized emittances at a higher injection energy, its extraction components and high-energy beam transport lines would not necessarily accommodate such larger normalized emittances without significant modifications. Arguments involving compatibility with existing machines on site favor preserving the output energy and bunch structure of the existing Booster. If the energy of the Phase-1 machine is lowered to 8 GeV, then an alternative upgrade scenario (to 4 MW of beam power for a muon collider) involving a second ring at 30 GeV suggests itself. That would serve the additional purpose of getting above Main Injector transition energy. Ultimately the choice may depend on whether potential future high energy physics programs that would directly use the output of the new facility will require beam energies greater than 8 GeV.

The proton bunch structure for the muon collider is specified at the pion production target; it might be achieved by combining several bunches at the target via chicanes. However, the present design adopts the straightforward approach of accelerating only two bunches. Regarding the bunch length, an rms value of 1 nsec is preferable, especially if it is desired to enhance the natural muon polarization, but the muon collider designers are willing to settle for 2 nsec. It is worth noting that the combination of high bunch intensities and short bunch lengths makes it difficult to avoid space-charge problems in the rings.

It might be wise to keep in mind the possibility of combining bunches at the production target, but the present design maintains the straightforward approach of accelerating the four bunches now required. An rms bunch length of 2 or 3 nsec would probably still suffice, but the present design aims for the more extreme value of 1 nsec. Shorter bunch lengths are preferred not only by those who would like to preserve the possibility of achieving higher muon polarizations for a neutrino factory or a muon collider, but also by those who would like to justify the choice of a higher beam energy for the proton driver. They argue that for a given longitudinal emittance, shorter bunches imply larger momentum spreads, and that for a given absolute momentum spread, the fractional momentum spread is larger at lower beam energy. However, even for beam energies as low as 8 GeV, the required dynamic aperture to contain the fractional momentum spread is not excessive.

The present design adopts the 15-Hz repetition rate of the existing proton source at Fermilab. The factor of three over the muon collider specification of 5 Hz can be

regarded either as a safety factor or as enabling operation of other physics programs at the same time as the muon collider.

The muon collider now needs the full repetition rate of 15 Hz to achieve its design luminosity because the process of production, capture, cooling, and acceleration of muons is less efficient than originally envisaged.

C.3. Synchrotron Design Concepts for Muon Production

The present plan for achieving the muon collider performance specifications calls for two rapid-cycling synchrotrons in series, each of which accelerates two bunches at a time. Table 1 presents output of a spreadsheet containing major parameters of the two rings. The strategy for achieving the required short bunches at the target while alleviating space-charge effects in the rings is to start with two relatively long bunches occupying most of the circumference of a small ring, and to do a bunch-narrowing rotation in longitudinal phase space just before extraction from each stage. In order to simplify matching of the bucket contour in the second ring to the rotated bunch distribution emerging from the first ring, the rf frequency in the second ring is chosen to be a multiple of that of the first ring. The rf frequency ratio, herein sometimes called the compression ratio, is chosen to be four in the present design.

After this was written, ESME simulations showed that it was possible to achieve the desired short rms bunch length of about 1 nsec with a single bunch rotation at extraction from the second ring, that is, without any rotation coming out of the first ring. Accordingly, the design now calls for the same rf frequency in the two rings.

The Laslett incoherent space-charge tune shift limits the beam brightness at low energy. The limitation on beam intensity can be raised by increasing the injection energy and by making the transverse emittances larger. Of course physical and dynamic apertures must be large enough to accommodate the large emittances. A useful approximation for the space-charge tune shift $\Delta \nu_{sc}$ at the center of a round Gaussian beam is

$$\Delta \nu_{sc} = -\frac{3 r_p N_{tot}}{2 \varepsilon_n \beta \gamma^2 B}$$

In this expression $r_p = 1.535 \times 10^{-18}$ m is the electromagnetic "radius" of the proton, N_{tot} is the total number of protons in the ring, ε_n is the 95% normalized transverse emittance, β and γ are the usual Lorentz kinematical factors, and B is the bunching factor, defined as the ratio of the average beam current to the peak current. Note that B is always less than or equal to one.

Both rings have peak dipole fields of 1.3 T in order to keep the ring circumferences relatively small while still allowing straightforward design of the conventional magnets. The magnet design is discussed elsewhere.

Even though the circumference of the larger ring was increased 50%, it was still necessary to raise the peak dipole field to 1.5 T.

Both rings employ separated-function lattices with flexible momentum compaction in order to raise transition above the extraction energy. This not only avoids accelerating through transition but also provides other advantages. Intense beams are not subject to certain instabilities such as the negative-mass instability below transition and empirically seem less susceptible to other instabilities such as microwave instability. Also, the negative natural chromaticity is beneficial for stabilizing the beam below transition, thereby perhaps obviating the need for sextupole correctors, especially in the first ring. Having transition not too far above extraction also provides substantial bucket area in which to accomplish beam-shortening rf manipulations. In the present design the transition energy is chosen to make the synchrotron frequency in the final stationary bucket high enough to accomplish the bunch rotation in less than about half a millisecond.

Careful design of the beam pipes for both rings is required in order to manage eddy-current effects. Two approaches are under consideration. One is a thin metal pipe with water cooling and eddy-current coil corrections integrated on the pipe as in the AGS Booster. The other is ceramic beam pipe with some sort of interior cage to carry beam image currents as in ISIS.

The present design calls for external vacuum skins for both rings as in the existing Fermilab Booster. However, unlike the Booster, some kind of conductive liner will be provided to carry beam image currents.

The first ring operates at a harmonic number h=2. This allows the two bunches to be formed directly and accelerated with efficient use of the whole circumference in order to keep the bunching factor large. An important paradoxical implication of the tune shift formula is that the tune shift for a given total number of particles is independent of the ring circumference; that is, a small ring will accommodate just as many particles as a large ring, other factors being equal. Space-charge effects are alleviated by a high injection energy (1 GeV kinetic) and large normalized transverse 95% emittances (200 π mm-mrad). (Large magnet apertures of order 15 cm are necessary to accommodate these emittances.) This produces a Laslett incoherent space-charge tune shift of 0.4 for a bunching factor of 0.25.

Of course the first ring now accelerates four bunches with the harmonic number h=4. The discussion of beam parameters required to make the space-charge tune shift tolerable uses the formula given above for Gaussian beams. The main body of this report assumes that injection painting can produce more favorable transverse and longitudinal form factors, resulting in a reduction of the tune shift by about a factor of two. That same factor of two could be used in this appendix to reduce the transverse emittances and/or the injection energy from the linac. The observation that both rings will require similar physical apertures is still true; the 200π beam emittance of the small ring is

compensated by the higher injection energy and smaller circumference, implying smaller beta functions. Hence magnet and rf cavity apertures can be similar if not identical.

The transfer energy of 4.5 GeV between the two rings is chosen to equalize the space-charge tune shift in the two rings. In the tune shift formula, there are two factors of γ. Roughly speaking, one factor of γ is used to make up for the larger circumference of the second ring; the other factor of γ is used to compensate for the shorter bunch length resulting from the bunch rotation. Both effects reduce the bunching factor in the second ring. (The bunching factor of 0.25 in Table 1 is the "bucket bunching factor"; it uses the average beam current over one rf wavelength rather than over the whole ring.)

Eliminating the bunch rotation upon extraction from the first ring makes the bunch longer when it is injected into the second ring. That means that a transfer energy of 3 GeV is high enough to equalize the space-charge tune shifts in the two rings, so the energy of the first ring was lowered to that value. At that energy it may not be necessary to employ flexible momentum compaction in order to avoid going through transition.

The design of the required 1-GeV linac is discussed elsewhere; here only a brief overview is given. H^- injection is used. It is assumed that the injected beam will be chopped and injected into pre-existing buckets to achieve high capture efficiency; the detailed optimization of this process is just beginning. (In simulations, adiabatic capture does not work well at the low rf frequencies considered here.) A debuncher is included to allow injection of small momentum spread, should this prove beneficial for creation of relatively small longitudinal emittance at injection into the first ring. The specified bunch rotations at extraction from each ring are expected to create momentum spreads of order 1% with longitudinal emittances of order 1 eV-sec per bunch. Such spreads would contribute a few centimeters in quadrature to the beam size for a short period before extraction from each machine. This is thought to be tolerable, given the large apertures required in any case.

As the first sentence implies, an energy upgrade to 1 GeV and an intensity upgrade to 10^{14} protons per cycle are described in TM-2021. The site layout depicted on the cover of this proton driver design study provides space for a 600-MeV "afterburner" downstream of the existing 400-MeV linac as well as for matching the 400-MeV beam into it and debunching downstream of it. The above paragraph implicitly contains another interesting design idea, namely that the horizontal aperture required for the bunch rotation at extraction time should not exceed what is required to contain the injected beam.

The detailed design of the rf systems is discussed elsewhere. Only an overview of some of the design choices is presented here. The magnet/power supply circuit for each ring is a 15-Hz resonant system like that of the existing Booster, with dipoles and quadrupoles electrically in series. (Of course this implies that the second ring can accelerate only one batch of two bunches at a time from the first ring in muon production mode.) Adding about 15% of second harmonic to the magnet ramp reduces the required peak accelerating voltage by about 25%, which is probably worth doing, especially for the

second ring with its large voltage requirement. Table 1 shows a few rf parameters such as accelerating voltages (in the absence of second harmonic) and rf frequencies. One of the advantages of a two-ring system is that the two rings divide the work of accelerating the beam. The rf system of the first ring is relatively modest because of its small circumference and small energy gain.

Table 1 is not reproduced here because the numbers are mostly obsolete. However the ideas expressed in this paragraph are still pertinent. Both rings would probably have Finemet cavities operating around 5 MHz to 7.5 MHz. The rf systems in the second ring would have to be modified to accommodate the higher beam power and beam loading of Phase II.

ESME simulations of longitudinal motion are underway; results to date have conformed qualitatively with expectations. In particular, the simulated rms bunch length at output of the second stage is consistent with a simple back-of-the-envelope estimate of 2 nsec, as desired. Accelerator studies at the Fermilab Booster and the Brookhaven AGS have begun to test bunch-narrowing concepts. Further work is in order, both experimental and computational, to optimize the bunch-shortening strategy. Also, longitudinal space-charge effects are important in these simulations; significant but tolerable emittance growth is predicted. High injection energies help to alleviate these longitudinal effects, which result from space-charge voltages having the same $1/\beta\gamma^2$ kinematic dependence as the transverse tune shifts. Incorporation of tunable inductive inserts in the rings is under consideration to compensate the space-charge voltages below transition. An experimental program is underway in collaboration with Los Alamos to study the effects of inductive inserts on the beam in the PSR.

Inductive inserts were installed in the PSR and provided the intended benefits. The accelerator studies at the Fermilab Booster and the Brookhaven AGS demonstrated the expected bunch-shortening. However, further study of bunch-shortening strategies is in order for the particular parameters of the second ring.

C.4. Meeting the Needs of the Rest of the Program

Within the general framework of multiple rings in series with bunch rotations before extraction from each stage, there is considerable flexibility in the choice of parameters (including the number of rings!). The parameters can be chosen in order to match the output beam to the needs of the rest of the physics program.

The present design starts with the choice of the circumference of the second ring to match that of the existing Booster. The output energy of about 16 GeV then results from an assumed dipole packing fraction of 0.575 and from the estimation that a dipole field of 1.3 T is about the highest reasonable choice that is consistent with straightforward design of magnets having thin silicon steel laminations. Driving such magnets into saturation would cause significant heating of the magnet yoke as well as potential problems with tracking of the dipoles and quadrupoles.

The issues limiting the maximum magnetic field were examined and found to allow a somewhat higher value of 1.5 T. Nevertheless, the circumference of the second ring had to be increased to 1.5 times that of the existing Booster in order to reach 16 GeV. The packing fraction of 0.575 (like that of the existing Booster) could not be achieved for a separated-function lattice with flexible momentum compaction and zero-dispersion straight sections.

The harmonic numbers of the two rings (2 and 21) and their respective circumferences are chosen in such a way that the bucket spacing in both is an integral multiple of the canonical Fermilab bucket spacing of 5.645 m. In particular, the circumference of the first ring is 8/21 times that of the second ring. Thus the bucket spacing in the second ring is four times, and that of the first ring is sixteen times, that of the Tevatron and Main Injector. The bunch structure resulting from either machine will then fall into buckets of any of the existing Fermilab rings. Thus the existing rf systems in downstream machines need not be replaced.

There are numerological constraints on the choices of circumferences and harmonic numbers to preserve compatibility between the two rings and with existing rings on site. The circumference of the small ring is chosen to be one-third that of the large ring, or in other words, half that of the existing Booster. That implies that the large ring should run at h=12 to match the bunch spacing of the small ring. Since the rf systems of the large ring must provide a large frequency swing anyway to accommodate injection at 400 MeV, it should also be possible to run them at h=12, instead of h=18, when the injection energy is 3 GeV.

An important design idea results from the realization that the Main Injector is grossly mismatched to the capabilities of the first ring. The Main Injector is incapable of handling the bunch intensities of 5×10^{13} that the muon collider requires, and the fill time would be excessive at the rate of two bunches every 66.6 msec. The normalized emittance of 200 π mm-mrad from the first ring, or 240 π mm-mrad from the second ring, also greatly exceeds the normalized acceptance of the Main Injector, which is specified as 40 π mm-mrad at 8 GeV and hence would be about 80 π mm-mrad at 16 GeV.

The mismatch between the first ring and the Main Injector can be circumvented by the simple expedient of bypassing the first ring when filling the Main Injector. All 21 buckets of the second ring would instead be filled directly from the linac using H⁻ injection. The Main Injector could then be filled with one or more Booster-length batches just as presently planned. However, since the rf bucket length of the second ring is four times that of the Main Injector, every fourth bucket of the Main Injector would contain beam.

The idea of bypassing the first ring is obviously incorporated in the main body of the design presented here since the first ring does not exist in Phase 1. Other changes are obvious; for example, there are 18 buckets in the second ring instead of 21, and every seventh Main Injector bucket would be filled.

It is worth noting that the first ring could be omitted if the muon collider does not materialize. However, if the first ring exists and the muon collider is not running, the output of the first ring could be used directly to support low-energy physics programs while the second ring is used to feed the Main Injector.

The capabilities of the second ring at an injection energy of 1 GeV are well-matched to those of the Main Injector, as can be seen by the following scaling. A normalized acceptance of 240 π mm-mrad at 4.5 GeV scales with momentum to 76 π mm-mrad at 1 GeV, closely matching the 80 π mm-mrad acceptance of the Main Injector at 16 GeV. The bunch intensity of 5×10^{13} at the space-charge limit at 4.5 GeV scales as $\beta^2 \gamma^3$ to 1.8×10^{12} per bunch, or 3.75×10^{13} per Booster-length batch, or 2.25×10^{14} per six Booster-length batches, at 1 GeV. This is 7.5 times the Main Injector design intensity. According to Weiren Chou, the Main Injector seems capable, with upgrades, of accelerating five times its design intensity. However, Chou's calculations were done for an injection energy of 8 GeV and with adjacent buckets filled, and they ought to be redone for the present case.

Redoing this scaling for the present parameters is left as an exercise for the reader.

The strategy of sometimes bypassing the first ring has several ramifications. It suggests a layout that has both rings tangent to the line from the linac as in Figure 1. The first ring should have a long straight section that supports both H⁻ injection and 4.5 GeV extraction. The second ring needs a long straight section that supports H⁻ injection at 1 GeV and proton injection at 4.5 GeV. Scaling from the Fermilab Booster implies that H⁻ injection at 1 GeV requires about a 12-meter straight section. This suggests a racetrack configuration for both rings.

Figure 1 is not reproduced here; the same features are contained in the picture on the cover of this report. Of course the shape of the second ring is now triangular, and the transfer energy is now 3 GeV rather than 4.5 GeV.

If it is desired to interleave cycles, some of which go into the first ring and some directly into the second, on a short time scale, then the linac must be capable of asynchronous operation. This requires further study.

The main parameters of Phase II, i.e., an upgrade to 4 MW from Phase I, are listed in Table C.1. As a comparison, the parameters of the present proton source and Proton Driver Phase I are also listed.

Table C.1. Proton Driver Parameters of Present, Phase I and Phase II

Parameters	Present	Phase I (MI, ν-Fact)	Phase II (μ-Coll)
Linac (operating at 15 Hz)			
Kinetic energy (MeV)	400	400	1000
Peak current (mA)	40	60	80
Pulse length (μs)	25	90	200
H⁻ per pulse	6.3×10^{12}	3.4×10^{13}	1×10^{14}
Average beam current (μA)	15	81	240
Beam power (kW)	6	32	240
Pre-Booster (operating at 15 Hz)			
Extraction kinetic energy (GeV)			3
Protons per bunch			$2,5 \times 10^{13}$
Number of bunches			4
Total number of protons			1×10^{14}
Normalized transverse emittance (mm-mrad)			$200 \, \pi$
Longitudinal emittance (eV-s)			2
RF frequency (MHz)			7.5
Average beam current (μA)			240
Target beam power (MW)			720
Booster (operating at 15 Hz)			
Extraction kinetic energy (GeV)	8	16	16
Protons per bunch	6×10^{10}	1.7×10^{12}	2.5×10^{13}
Number of bunches	84	18	4
Total number of protons	5×10^{12}	3×10^{13}	1×10^{14}
Normalized transverse emittance (mm-mrad)	$15 \, \pi$	$60 \, \pi$	$200 \, \pi$
Longitudinal emittance (eV-s)	0.1	0.4	2
RF frequency (MHz)	53	7.5	7.5
Extracted bunch length σ_t (ns)	0.2	1	1
Average beam current (μA)	12	72	240
Target beam power (MW)	0.1	1.2	4